

# SPECIFICATION

Electronic Version 1.2.8

Stylesheet Version 1.0

## ***INFORMATION RETRIEVAL FROM A COLLECTION OF INFORMATION OBJECTS TAGGED WITH HIERARCHICAL KEYWORDS***

### **Background of Invention**

- [0001] *FIELD OF THE INVENTION* This invention relates in general to methods and data processing system readable media, and more particularly, to data processing system-implemented methods of formulating queries and searching for a plurality of information objects and data processing system readable media having software code for carrying out those methods.
- [0002] *DESCRIPTION OF THE RELATED ART* A goal of information retrieval systems is to allow efficient access to selected documents or other kinds of information objects from a repository. The user of such a system may be interested in knowing the existence and location of the available information objects that are relevant to a specific request or query.
- [0003] A common approach used in information retrieval systems is to associate one or more keywords with each information object. The set of all known keywords comprises the "master set" of keywords. To form a query, the user provides one or more keywords, which may or may not be drawn from the master set. The information retrieval system then returns each information object for which one or more of its associated keywords matches one or more of the keywords in the query. As a further step, a mathematical formula can be applied to the number of keyword matches to provide a scalar that is associated with each information object returned by the query. The scalar serves as a "relevance score" that indicates the degree to which the particular information object matches the query. This approach can be generally termed "keyword-matching" and there are many specific embodiments used in practice. Some difficulties with the keyword-matching approach are set forth in the following paragraphs.

- [0004] First, the user of the system may not know or be able to grasp all of the possible keywords in the master set. In this case, the user may provide queries that contain keywords that are not used in the master set. This reduces the effectiveness of the system, particularly when the master set includes keywords that have closely related meanings in a particular application, and a simple match cannot make use of this information. For example, assume the repository contains documents describing fruits and vegetables, and a treatise on tomatoes has been assigned the keyword "nightshade" because it also includes discussions of eggplant and potatoes. The user desiring information on tomatoes might enter a query such as "tomatoes" and this query would fail to match the treatise on the nightshade family, even though that document is relevant to the user's purpose.
- [0005] Second, the mathematical formulae that are widely described and used to compute relevance scores may not take advantage of the relationships among keywords that are inherent in any specific information repository. For example, given a repository that contains documents on fruits and vegetables, systems that compute a relevance score based only on the number of keyword matches have no way to incorporate the fact that a document tagged with keywords "nightshade" and "treatise" should more closely match the query pair "tomato" and "treatise" than the query pair "lamp" and "treatise." Attempts to address these shortcomings have been proposed, but the methods fail to fully address the problems users may encounter. Some systems have been developed that organize the keywords into a hierarchical tree structure. This, by itself, is not a solution, as will become evident in some of the paragraphs that follow.
- [0006] A system described in U.S. Patent 6,094,652 ("Faisal") places keywords into a hierarchical structure. The hierarchy expresses the associations among the keywords in the repository. When responding to a user query, the system suggests keywords from the hierarchy that broaden or narrow the scope. The system also suggests keywords that represent concepts that are neither broader nor narrower but are related by means of an explicit cross-link among the nodes in the keyword hierarchy. The user can refine his or her query in an interactive and iterative fashion.
- [0007] A system described in U.S. Patent 6,098,066 ("Snow") arranged the information objects into a document hierarchy (a tree data structure). Each node of the hierarchy corresponds to a category and contains at least one document. The user of the system has the option of restricting their search to the documents branching from a specific category (which these

authors term a "directed" search) or searching all documents in the repository (which these authors term an "undirected" search). The user may restrict the number of documents returned by the system by focusing on a particular category, while leaving the user with the option of searching the entire repository if desired.

- [0008] A system described in U.S. Patent 5,991,756 ("Wu") places documents into a hierarchical structure. The system retrieves documents that match one or more query keywords directly or match "indirectly" by being located as a child node to a document in the document hierarchy that matches directly one or more of the query terms.
- [0009] A system described in U.S. Patent 5,630,125 ("Zellweger") places documents into a hierarchical structure that has one or more paths leading to a given document. The system provides an interactive method that allows the user to formulate a final query by navigating the hierarchy structure to the desired documents. Multiple paths support synonyms and allow the user to clarify word meaning in a given context.
- [0010] A system described in U.S. Patent 5,787,417 ("Hargrove") is highly similar to that described by Zellweger in that it provides an interface for allowing the user to interactively navigate the hierarchy of the repository to locate the desired information objects.
- [0011] A textbook by C.J. Van Rijsbergen (Information Retrieval, 2<sup>nd</sup> Ed) describes a general strategy for information retrieval by keyword matching. It also gives the mathematical formulae that can be used to transform the combination of a "query vector" and a "document vector" into a final "relevance score" that can be used to rank the documents returned by a retrieval system according to their degree of relevance to the query.
- [0012] Each of the systems in those documents has at least one limitation or disadvantage in some applications.
- [0013] Systems that require the user to interactively refine their query (such as those described by Faisal, Zellweger, and Hargrove) are inherently more time consuming for the user than a system that returns results in response to a single query. Further, human interfacing with a computer costs a company valuable human resources. In some applications (such as those described in the next section), the information retrieval is automated, and there is no opportunity to refine or otherwise change the query before searching begins.

- [0014] Systems that restrict the retrieved documents to those with a particular ancestry in a document hierarchical structure (such as those described by Faisal, Snow, and Wu) can fail to return relevant documents outside their hierarchical search path unless there have been many cross-links provided (such as in the system described by Faisal). Cross links must be created and maintained manually, a time-consuming and error-prone process.
- [0015] Several of the prior systems do not prescribe a method for assigning a relevance score between the query and the documents in the repository (such as the systems described by Zellweger and Hargrove). It is often convenient for the users to have a relevance score to help them estimate their level of interest in the returned documents. Furthermore, systems that restrict the search path to a particular set of child nodes in the hierarchy (such as that described by Wu) cannot provide relevance scores for documents that lie outside the restricted set of child nodes. In some applications, this means that not all documents can be assigned a relevance score in response to a given query.

## Summary of Invention

- [0016] Embodiments of the present invention do not suffer from the problems seen with prior art methods and systems. A user is not required to interactively refine a search because a data processing system can be programmed to automatically expand an original query having original keywords to an expanded query that includes friend keywords of the original keywords. Searching may be performed that cover parts of one or more hierarchies because keywords outside a specific ancestry may be used. Further, the documents are not required to be placed within a document hierarchy. A cross-link system is not required, which saves valuable money and human resources.
- [0017] In one set of embodiments, a data processing system-implemented method of searching for a plurality of information objects can comprise receiving a first signal that includes or is used to form a first query. The first query may include a first keyword within a hierarchy. The method can also comprise expanding the first query to a second query. The second query may include the first keyword and a second keyword within the hierarchy. The method can further comprise searching the database using the second query and finding a first identifier for a first information object that corresponds to the second query.
- [0018] In another set of embodiments, a data processing system-implemented method of

formulating a query can comprise receiving a first signal that includes or is used to form a first query having a first keyword. The method can also comprise automatically expanding the first query to a second query. The second query may include the first keyword and a second keyword that is not present within the first query.

- [0019] In still other embodiments, a data processing system readable medium can have code embodied within it. The code can include instructions executable by a data processing system. The instructions may be configured to cause the data processing system to perform the methods described herein.
- [0020] The foregoing general description and the following detailed description are exemplary and explanatory only and are not restrictive of the invention, as defined in the appended claims.

## Brief Description of Drawings

- [0021] The accompanying drawings, which are incorporated in and constitute a part of this specification, illustrate embodiments of the invention and together with the description, serve to explain the principles on of the invention.
- [0022] FIG. 1 includes an illustration of a hardware architecture for carrying out methods of searching a database;
- [0023] FIG. 2 includes an illustration of a data processing system storage medium including software code;
- [0024] FIG. 3 includes an illustration of a hierarchy of keywords;
- [0025] FIG. 4 includes a flow diagram for adding keywords, friends, and information objects to a database; and
- [0026] FIG. 5 includes a flow diagram for obtaining information objects related to a keyword and its friends.
- [0027] Skilled artisans appreciate that elements in the figures are illustrated for simplicity and clarity and have not necessarily been drawn to scale. For example, the dimensions of some of the elements in the figures may be exaggerated relative to other elements to help to improve understanding of embodiments of the present invention.

## Detailed Description

- [0028] Reference is now made in detail to the exemplary embodiments of the invention, examples of which are illustrated in the accompanying drawings. Wherever possible, the same reference numbers will be used throughout the drawings to refer to the same or like parts (elements).
- [0029] The present invention can include a data processing system-implemented method or a data processing system readable media having software code for carrying out the method. The method can comprise formulating queries, searching for a plurality of information objects, or a combination thereof. In a specific embodiment, an original query with at least one keyword can be automatically expanded to an expanded query that includes at least one keyword that is not in the original query. The expanded query may be used to search for information objects that are relevant to the expanded query.
- [0030] Before discussing embodiments of the present invention, a hardware architecture for using embodiments is described. FIG. 1 illustrates an exemplary architecture and includes a client computer 12 that is bi-directionally coupled to a network 14, and a server computer 16 that is bi-directionally coupled to the network 14 and database 18. The client computer 12 includes a central processing unit ("CPU") 120, a read-only memory ("ROM") 122, a random access memory ("RAM") 124, a hard drive ("HD") or storage memory 126, and input/output device(s) ("I/O") 128. The I/O devices 128 can include a keyboard, monitor, printer, electronic pointing device (e.g., mouse, trackball, etc.), or the like. The server computer 16 can include a CPU 160, ROM 162, RAM 164, HD 166, and I/O 168.
- [0031] Each of the client computer 12 and the server computer 16 are examples of data processing systems. ROM 122 and 162, RAM 124 and 164, HD 126 and 166, and the database 10 include media that can be read by the CPU 120 or 160. Therefore, each of these types of memories includes a data processing system readable medium. These memories may be internal or external to the computers 12 and 14.
- [0032] The methods described herein may be implemented in suitable software code that may reside within ROM 122 or 162, RAM 124 or 164, or HD 126 or 166. In addition to those types of memories, the instructions in an embodiment of the present invention may be contained on a data storage device with a different data processing system readable storage medium, such as a floppy diskette. FIG. 2 illustrates a combination of software code elements 204, 206, and 208 that are embodied within a data processing system readable medium 202, on a floppy diskette

200. Alternatively, the instructions may be stored as software code elements on a DASD array, magnetic tape, conventional hard disk drive, electronic read-only memory, optical storage device, CD ROM or other appropriate data processing system readable medium or storage device.

[0033] In an illustrative embodiment of the invention, the computer-executable instructions may be lines of compiled C<sup>++</sup>, Java, or other language code. Other architectures may be used. For example, the functions of the client computer 12 may be incorporated into the server computer 16, and vice versa. FIGs. 4 and 5 include illustrations, in the form of flowcharts, of the structures and operations of such a software program.

[0034] Communications between the client computer 12 and the server computer 16 can be accomplished using electronic or optical signals. When a user (human) is at the client computer 12, the client computer 12 may convert the signals to a human understandable form when sending a communication to the user and may convert input from a human to appropriate electronic or optical signals to be used by the client computer 12 or the server computer 16.

[0035] Attention is now directed to data preparation and system initialization for searching. During data preparation, a master list of keywords (referred to as the Master Keyword List) is generated and arranged into one or more sets of hierarchical relationships or "trees." For example, FIG. 3 includes a tree (hierarchy) with 12 keywords: *B*, *C*, ..., *M*. The root of the tree is the node *A* 302. Nodes *B*312, *C*314, and *D*316 are the children of node *A* 302. Nodes 321 *E*321, *F*322, and *G*323 are the children of node *B*312. Nodes *H*326, *I*327, and *J*328 are the children of node *D*316. Nodes 332 *K*332, *L*334, and *M*336 are the children of node *I*327. Nodes *C*314, *E*321, *F*322, *G*323, *H*326, *J*328, *K*332, *L*334, and *M*336 have no children.

[0036] The dashed lines 310, 320, and 330 are used as a point of reference to divide "generations" of nodes. Each node belongs to a specific "generation" that is equal to the number of ancestor nodes between the given node and the root of the tree. For example, nodes 332, 334, and 336 are members of the third generation (generation = 3) because these nodes have three ancestors (nodes 327, 316, and 302). The significance of the generations will become apparent later. Note that FIG. 3 can allow for the identification of the "lowest common ancestor" node. For example, when comparing nodes 332 and 336, the lowest common ancestor node is node 327, which is the parent node for each of nodes 332 and 336. When comparing nodes 326 and 336, the lowest common ancestor node is node 316, which is the parent node to node 326 and the

grandparent node of node 336.

- [0037] The name of the tree can be the root, which in this example can be "tree A." If tree A were the only tree, then { B , C , ... , M } may also be the master list of keywords. Each tree can be represented as a relational database table, as shown for this example in Table I. Additional tables may be present for other keyword hierarchies.

- [0038] *Table I: Keyword Hierarchy Table*

- [0039]

Keyword / Node	Parent	Generation
B	A	1
C	A	1
D	A	1
E	B	2
F	B	2
G	B	2
H	D	2
I	D	2
J	D	2
K	I	3
L	I	3
M	I	3

- [0040] Referring to FIG. 4, the keyword hierarchy is produced (circle 412) and stored at part of the Keyword Hierarchy Table 414 that may be within database 18.

- [0041] The database 18 may include a repository of information objects. The information objects themselves may include documents, products, electronic discussion archives, code fragments, and any other computer representations of knowledge or information. Each information object can have a unique identifier, hereafter called the "object ID." In addition, each information object may have other important properties, such as its language, access control parameters, object type (document, software product, etc.), and the like.

- [0042] A file or database table can specify a set of keywords relevant to each information object, which will be referred to as the "Info Object Keyword Table." All keywords related to the information objects should be members of the Master Keyword List. In other words, keywords

are related to information objects (circle 422) and can be stored as part of the Info Object-Keyword Table 424 as seen in FIG. 4. In one embodiment, a relevance rating for each keyword may be provided and can represent the degree of relevance between a keyword and an information object. The relevance rating can be assigned by subject matter experts who assign keywords to information objects and populate the information object repository. This data can be contained in a relational database table where each row contains an object ID, a keyword (or keyword identifier), and a relevance rating (e.g., from 1 to 10), as shown for example in Table II. An object ID can appear multiple times in this table when more than one keyword is considered relevant to the information object.

[0043] Table II: Relevance Between Object IDs and Keywords.

[0044]

Object Id	Keyword	Rating
obj1	C	8
obj1	F	7
	B	5
obj2	H	9
	I	9
obj3	E	10
obj4	J	9
obj4	M	6
obj5	F	7
obj5	J	4
obj5	L	8
obj6	D	6
obj6	C	8
obj7	H	9
obj8	B	4
obj8	D	8

[0045] Note that the relevance rating assigned to a given information object-keyword pair need not be the same as the rating assigned to a different information object-keyword pair, even if the keyword is the same in both cases. In Table II, for example, the keyword J is relevant to information object "obj4" with relevance rating of 9, and to "obj5" with relevance rating of 4. This means that keyword J is more relevant to obj4 than to obj5.

- [0046] The Keyword Hierarchy Tables may be used to produce a set of "friend" keywords (and corresponding association scores between a keyword-friend pair) for every keyword in the Master Keyword List. After a keyword-friend association score can be calculated (circle 432), and that information may be stored in the Keyword-Friend Table 434 of database 18.
- [0047] Attention is now directed to some of the details in determining keyword-friend association scores. A number of different methods can be used to determine the association scores between keywords and their friends. In one embodiment, the association score may be determined using a tree distance algorithm, further described below. The pairing of keywords with their friends can be maintained in the relational database table 424 (within database 18, for example), with an entry for each keyword / friend pair. Every keyword is a friend of itself, with the maximum possible association score. A keyword's other friends can be given by further entries in the Keyword-Friend Table 434, one for each distinct pair of keywords, along with the association score for that pair. An example consistent with FIG. 3 may include association scores that range from 1 to 10. As shown in Table III.
- [0048]
- [0049]
- [0050]
- [0051]
- [0052]
- [0053]
- [0054] Table III: Keyword Friend Table
- [0055]

Keyword	Friend	Distance	Score
B	B	0	10
B	E	2	8
B	F	2	8
B	G	2	8
C	C	0	10
C	B	6	4
C	D	6	4
D	D	0	10
D	H	2	8
D	I	2	8
D	J	2	8
E	E	0	10
E	B	2	8
F	F	0	10
F	B	2	8
G	G	0	10
G	B	2	8
H	H	0	10
H	D	2	8
I	I	0	10
I	K	1	9
I	L	1	9
I	M	1	9
J	J	0	10
J	D	2	8
K	K	0	10
K	I	1	9
L	L	0	10
L	I	1	9
M	M	0	10
M	I	1	9

[0056]

[0057] Associated with each link in the keyword hierarchy (represented by the arrows in FIG. 3) is a weight. The weight is equal to the highest generation number in the tree minus the generation of the parent node in the link. For example, the links between node /327 and its children *K* 332, *L* 334, and *M* 336 have a weight equal to 1, which is the highest generation number in the tree (3) minus the generation of the parent node /(2). The dotted horizontal lines in FIG. 3 indicate the generations, and line 310 can correspond to a weight of "3," line 320 can correspond to a weight of "2," and line 330 can correspond to a weight of "1."

[0058] The association score between any two keywords in the hierarchy may be determined in two acts:

[0059] 1. compute the "tree distance" between the two keywords, then

[0060] 2. transform the tree distance according to a mathematical equation to get the final association score.

[0061] To compute the tree distance, the method can use the following algorithm. Trace the ancestry of each node up to the lowest common ancestor. For each link that is used to get to the lowest common ancestor, maintain a sum of the weights. One embodiment may use a sum of the weights squared. (In general, a user can define the tree distance to be the sum of any bias function applied to the weights; power-laws may be particularly useful.)

[0062] For example, let  $d(x,y)$  denote the tree distance between nodes  $x$  and  $y$ , where  $x$  and  $y$  are nodes in the tree. To compute  $d(L,H)$ , a path can start at node  $L$  334 and traces across dotted line 330 to its parent node  $I$  327, and then across dotted line 320 to its grandparent node  $D$  316. Node  $D$  316 is the lowest common ancestor between nodes  $L$  334 and  $H$  326. Starting at node  $H$  326, a path can reach node  $D$  316 by crossing line 312. Thus, the distance can be calculated as  $(1 + 2)$  for going from node  $L$  328 to node  $D$  316, and adding 2 for the distance from node  $H$  326 to  $D$  316. In more explicit algebraic notation,

[0063]  $d(L,H) = d(L,I) + d(I,H)$

[0064]  $d(L,H) = d(L,I) + (d(I,D) + d(H,D))$

[0065]  $d(L,H) = 1 + 2 + 2$

[0066]  $d(L,H) = 5.$

[0067] The following are some more examples:

[0068]  $d(L,M) = 1 + 1 = 2$

[0069]  $d(L,H) = (1 + 2) + 2 = 5$

[0070]  $d(L,C) = (1 + 2 + 3) + 3 = 9$

[0071]  $d(H,C) = (2 + 3) + 3 = 8$

- [0072]  $d(E,C) = (2 + 3) + 3 = 8$
- [0073]  $d(H,E) = (2 + 3) + (2 + 3) = 10$
- [0074]  $d(L,E) = (1 + 2 + 3) + (2 + 3) = 11$
- [0075] Other methods may be used to determine the tree distance. Between sibling nodes (child nodes from a common parent node), a symmetric distance matrix may be generated to determine scale of distances between those sibling nodes. That is, the distance between any two children of a parent node can be determined by multiplying appropriate entry in the distance matrix by the sum-of-weights distance.

- [0076] For example, in FIG. 3, the designers of the tree may choose to define a child distance matrix for node *D* 316 as the following:

[0077]

	H	I	J
H	1	1	1.5
I	1	1	1
J	1.5	1	1

- [0078] This matrix has two properties. First, the diagonal entries are all equal to one. Second, it is symmetric. The elements of the matrix can be denoted by  $M(x,y)$ , so that (for example)  $M(H,J) = 1.5$ .
- [0079] A method for calculating an association score between nodes can use this matrix to scale the sum-of-weights distance between the child nodes of the lowest common ancestor. Using only the sum-of-weights distance on the tree in FIG. 3 would produce the result  $d(L,H) = d(L,J) = 5$ . The child-distance matrix allows the designers of the tree to express a closer relationship between some children than between others. In this example, nodes *H* 326 and *I* 327 are more closely related to each other than to node *J* 328 even though all three share the same parent.
- [0080] In one example, the method can use the child-distance matrix to scale the distance between children of node *D* to express the closer relationship between nodes *H* 326 and *I* 327:

[0081]  $d(L,H) = d(L,I) + M(I,H)*d(I,H)$

[0082]  $d(L,H) = d(L,I) + M(I,H)*(d(I,D) + d(H,D))$

[0083]  $d(L,H) = 1 + 1*(2 + 2)$

[0084]  $d(L,H) = 5.$

[0085] Between nodes J 328 and L 334, the calculation may be:

[0086]  $d(L,J) = d(L,I) + M(I,J)*d(I,J)$

[0087]  $d(L,J) = d(L,I) + M(I,J)*(d(I,D) + d(J,D))$

[0088]  $d(L,J) = 1 + 1*(2 + 2)$

[0089]  $d(L,J) = 5.$

[0090] While the matrix may usually be symmetric, symmetry is not required.

[0091] This procedure can achieve a desirable effect, in that it gives designers additional flexibility to define quantitative relationships among the keywords in the hierarchy, and these relationships can be used to provide superior information retrieval results. The association scores may be determined automatically by server computer 16 based at least in part upon positions of a keyword and its friend Keyword within the hierarchy.

[0092] The method can produce an association score between two keywords by computing a distance between the keywords (using the sum-of-weights or the sum-of-weights plus child-distance matrix method) and then applying a transform to give the highest association scores to those keywords with the lowest values of the tree distance.

[0093] One embodiment can use the following transformation. Let C denote the maximum desired association score. The association score between any object and itself is equal to C. Let d denote the tree distance. Let  $f(d)$  denote a monotonically increasing function of the argument d, and  $\text{int}(f(d))$  denote an integer value of  $f(d)$ . Let s denote an association score corresponding to tree distance d, and can be given by:

[0094]  $s = \max(0, c - \text{int}(f(d)))$

[0095] In this formula, the value of  $\max(x,y)$  is the greater of the arguments x and y and  $\text{int}(x)$  is the integer part of the argument x. The value "0" in the equation may be used so that s cannot

be negative. In one embodiment, the squares of the weights can be summed to obtain the tree distance d,  $f(d) = d^{**0.5}$  and C=10.

- [0096] To determine the "friend" keywords for a given keyword, the method can be used to compute the association scores between the given keyword and all other keywords in the hierarchy. The top N of the keywords with the greatest association scores become the set of friend keywords that are stored in a relational database table within database 18. One embodiment can use the keywords with the top 10 association scores to expand the query keywords. A user or code in software or hardware can be set the value of N.
- [0097] At this point in the process, data preparation and system initialization has been completed. The appropriate information may be stored with database 18 or other storage device having persistent memory.
- [0098] Query processing can now be performed. It is in this second phase, query processing, that information retrieval actually occurs. A query can be a set of keywords (or keyword IDs), generated by some specific end-user activity for some particular application. The nature of such applications and specific examples are discussed below. The query keywords are members of the Master Keyword List. By using a limited number of keywords from the Master Keyword List, searching can be performed faster compared to free-form searching.
- [0099] FIG. 5 includes a flow diagram of acts that can be performed when processing a query. Note that some of the acts may be optional and not required for all implementations. The method can comprise receiving from a client computer 12, an original (first) query that includes a first keyword within a hierarchy (block 502). After the original query is received, the server computer 16 can retrieve all the friend keywords (and association scores) for each keyword in the original query. A second keyword from the Keyword-Friend Table 424 can be identified as a friend of the first keyword, although the second keyword may not have been a keyword within the first query.
- [0100] The method further includes automatically expanding the original query to an expanded (second) query that includes the first keyword and the second keyword within the hierarchy (block 522). In this specific example, note that a third keyword may be present within the Master Keyword List but is not part of the expanded query. The third keyword may not be listed as a friend of the first keyword, or the third keyword may not have had a sufficiently high

enough association score compared to other friend keywords of the first keyword.

- [0101] The complete set of keywords from the original query and their corresponding friend keywords can form the expanded query.
- [0102] The association scores for the keyword-friend pairs may be used to determine which friends to use but can also be used in relevance scoring that will be described in more detail later.
- [0103] Below is an example using the hierarchy in FIG. 3 and the association scores in Table III.

[0104]

Original Query keywords:	C	K			
Expanded Query keywords:	C	B	D	K	I
Association scores:	10	4	4	10	9

- [0105] Keywords C and K can be examples of the first keywords, keywords B, D, and I can be examples of the second keywords, and keywords E, F, G, H, J, L, and M can be examples of the third keywords, which are not part of the expanded query.
- [0106] After the expanded query has been generated, the method can include searching the database 18 using the expanded query in block 542 of FIG. 5. The method also can include finding identifiers for information objects that correspond to the expanded query (block 544). More specifically, the keywords in the expanded query can be used to identify a set of relevant information objects via the Info Object Keyword Table 424. Only information objects with at least one keyword that is present in the expanded query may be considered relevant to the query.
- [0107] Table IV. Object identifiers, keywords, and relevance ratings.
- [0108]

Object Id	Keyword	Relevancy Rating
obj1	C	8
obj1	F	7
obj2	B	5
obj2	J	9
obj2	I	9
obj6	D	6
obj6	C	8
obj8	B	4
obj8	D	8

[0109] Table IV includes an exemplary set of relevant information objects corresponding to the expanded query. The relevance rating can be a rating of how relevant a keyword is to a specific information object. Note that obj8 is relevant to the expanded query but not to the original query because obj 8 only includes keywords *B* and *D* (second keywords) as relevant keywords.

[0110] Next, the method can calculate a relevance score for each of the identified information objects (block 562). There are many possible formulae for this calculation. In many instances, a weighted vector may be used for the relevance score. The weighted vector can be a product between two vectors of dimensionality D, where D is the total number of distinct keywords in the master keyword list, and each vector element is the relevance score for that keyword or zero if the keyword is absent. Users may find it useful to choose a formula which includes appropriate normalization to account for variable parameters that should not affect the final score spuriously, such as the number of keywords for a given information object. This consideration suggests the following formula for the relevance score R:

[0111]

$$R = \frac{\sum_{i=1}^M k_i q_i}{K}$$

[0112] where:

[0113] {K} is the set of keywords associated with the information object;

[0114] {k} is the set of relevance ratings for the information object keywords in {K};

- [0115]  $K = |\{K\}|$  is the number of keywords associated with the information object;
- [0116]  $\{Q\}$  is the set of keywords in the expanded query;
- [0117]  $\{q\}$  is the set of association scores for the expanded query keywords in  $\{Q\}$ ;
- [0118]  $\{M\}$  is the set of keywords from the intersection of sets  $\{K\}$  and  $\{K\}$ ;
- [0119]  $M = |\{M\}|$  is the number keywords in the set  $\{M\}$  (i.e. the number of matches);
- [0120]  $\{k_i\}, i=1,2,\dots,M$  is the subset of  $\{k\}$  corresponding to the elements of  $\{K\}$  in  $\{M\}$ ; and
- [0121]  $\{q_i\}, i=1,2,\dots,M$  is the subset of  $\{q\}$  corresponding to the elements of  $\{Q\}$  in  $\{M\}$ .
- [0122] The relevance score is determined as follows: find the sum, over each expanded query keyword that matches an information object keyword, of the product of the keyword's association score for the query and its relevance rating for the information object; and divide the sum by the number of keywords associated with the information object. This last act provides appropriate normalization to avoid arbitrarily enhancing the score of information objects that have a large number of keywords. Note that it is not necessary to normalize by the number of query keywords, since this is a constant for a given query. Note that the equation give above is not the only way to determine a relevance score, and therefore, should not be construed as a limiting. Table V includes the relevance scores obtained in this fashion for the given set of eligible information objects in Table IV and the expanded query.

[0123] Table V: Relevance Score Table.

[0124]

Object Id	Relevance Score
obj1	40
obj2	33.7
obj6	52
obj8	24

[0125] The method may further include sorting the identified information objects based on the relevance scores (block 564). The list of eligible information objects can be sorted from highest to lowest relevance score. After sorting, the method may send the sorted information object to

the client computer (block 582). The resulting list should provide the object IDs of all relevant information objects for the expanded query, in order of relevance, based on the original query. Although the calculating of the relevance score and sorting the identified information objects is optional, the information objects, sorted by relevancy score aid the user at the client computer by indicating the degree of relevance based on relevancy score. The final list for the example of Table V can be given by Table VI.

[0126] Table VI: Final List of Sorted Information Objects.

[0127]

Object Id	Relevance Score
obj6	52
obj1	40
obj2	33.7
obj8	24

[0128] The server computer 16 may send the client computer 12 the list seen in Table VI or a derivative of it. The information object IDs may include an alpha-numeric representation, a catalog number, or be replaced by a title, or even the information object itself (or the first few words of it) when information is seen by a user at client computer 16. Therefore sending an information object ID should be construed as including any one or more of the pieces of information listed in this paragraph.

[0129] Filtering acts may be used as an optimal part of the method. Filtering may select information objects by language, security level, length of document, or the like. The user may define the filtering criterion. The filtering criterion may be sent from the client computer 12 to the server computer 16 with the original query.

[0130]

Other specific embodiments are presented to illustrate some of the other features of the keyword-friend query method and system. One embodiment may serve as part of a suite of information retrieval systems for a corporate knowledge management system. The role of the method in this system is to provide employees, business partners, and customers of the corporation with efficient access to information objects that are relevant to a particular topic or user query. Access to the method and system can be through a software application made

available on the World Wide Web computer network via the HTTP communications protocol.

- [0131] In this embodiment there may be three hierarchies: one hierarchy may correspond to the subject matter category addressed by the information object ("category hierarchy"), another hierarchy may correspond to the project phase addressed by the information object ("phase hierarchy"), and still another hierarchy may correspond to the role, or intended audience, of the information object ("role hierarchy"). Each hierarchy is described by a database table that lists the unique identifiers of the keywords in that hierarchy, along with the unique identifiers of the parent keywords. The keyword corresponding to the root of the tree may have no parent keyword. The translation from unique keyword identifier to the keyword text can be provided by means of another database table.
- [0132] A further aspect of this embodiment can be a "keyword-type" weight value assigned to each of the keyword hierarchies. This can allow some types of keywords to be more significant than others, for example, a keyword of type "category" can be given a weight of 2 while "role" and "phase" keywords are given a weight of 1. All of the keywords in a given hierarchy can have an identical keyword-type weight. This weight value may be separate and in addition to the association score of a specific keyword for a specific information object, and can be used as an additional multiplicative factor for each term in the summation used in the relevance score.
- [0133] In one application of this embodiment, a user at client computer 12 may formulate a query by selecting one or more keywords from the three available hierarchies. The keywords can be displayed on a screen and selected by selection boxes displayed on the user's HTML browser. The keywords he or she selects can then be submitted to the server computer 16 that can search database 18 for the relevant information objects. The relevant information objects that are returned can then be sent from the server computer 16 and received by the client computer 12 where the user can see the results.
- [0134] In another application of this embodiment, the query may be formulated automatically according to the user's context on the Web site. By means of a series of questions presented to the user as they navigate the Web site, a query can be formulated based upon the question responses. For example, a series of question may reveal the user's interest in a particular subject matter category, and the keyword corresponding to that category can form the query that is submitted for processing. The relevant information objects that are returned are then presented to the user. Alternatively, the actions of the user at the website can be tracked and

information objects be presented to the user without having to ask any questions or receive a query from the user.

- [0135] Embodiments of the present invention have advantages over the prior art in the field of information retrieval. The methods can make use of the relationships among the keywords associated to each information object. These relationships can be expressed in the hierarchies to which the keywords belong. By taking advantage of this information, the methods can be used to find information objects in the repository that are close, but not exact, matches to an original query.
- [0136] The embodiments do not require interactive participation from the user beyond the specification of the original query. The invention may automatically expand the scope of the original query to include keywords with related meanings, so that matches can occur even on keywords the user did not think to enter in the original query.
- [0137] The methods may require only a few operations to compute a relevance score for an information object. Therefore, the methods can be performed on a data processing system in a time efficient manner.
- [0138] A further advantage seen with embodiments of the present invention is the ability to create derivatives to address other problems or to be used in other fields. Direct mail, electronic mail, and the World Wide Web provide marketers with an opportunity to target product offers to specific customers. These offers can include product recommendations. The process of constructing a product recommendation can be analogous to the process of information retrieval. Therefore, the methods can be used to provide product recommendations. The products (or more specifically unique product codes) themselves can serve the role of keywords. The product hierarchy maintained by many retailers can be used to establish the relationships among the keywords. The form of a query may be a set of products in which a customer has expressed interest (through purchase, request for information, etc.). At least one of the methods previously described can be used to return a list of related information objects (which are product identifiers) to the customer as a kind of product recommendation. Such recommendation may be valid under an assumption that consumers will be interested in products similar to those that they have purchased or browsed in the past. Many retailers maintain product hierarchies that express the similarity of items that are closely related in a tree.

[0139] In the foregoing specification, the invention has been described with reference to specific embodiments. However, one of ordinary skill in the art appreciates that various modifications and changes can be made without departing from the scope of the present invention as set forth in the claims below. Accordingly, the specification and figures are to be regarded in an illustrative rather than a restrictive sense, and all such modifications are intended to be included within the scope of present invention.

[0140] Benefits, other advantages, and solutions to problems have been described above with regard to specific embodiments. However, the benefits, advantages, solutions to problems, and any element(s) that may cause any benefit, advantage, or solution to occur or become more pronounced are not to be construed as a critical, required, or essential feature or element of any or all the claims. As used herein, the terms "comprises," "comprising," or any other variation thereof, are intended to cover a non-exclusive inclusion, such that a process, method, article, or apparatus that comprises a list of elements does not include only those elements but may include other elements not expressly listed or inherent to such process, method, article, or apparatus.